

File Views

Standard:

Why do we have file views? They tell MPI what I/O will occur in the future

<https://www.mpi-forum.org/docs/mpi-4.1/mpi41-report.pdf>

File views

Intro p.638 [definition of file view concept]

“Initially, all processes view the file as a linear byte stream, and each process views data in its own native representation (no data representation conversion is performed). (POSIX files are linear byte streams in the native representation.) The file view can be changed via the MPI_FILE_SET_VIEW routine.”

File_open p. 640

File_set_view / File_get_view pp. 650-653

File_read_at_all pp. 657-658

File_write_at_all p. 659

File_iread_at_all p. 661

File_iwrite_at_all p. 663

File_read_all pp. 665-666

File_write_all pp. 667-668

File_iread_all pp. 669-670

File_iwrite_all p. 671

Collective I/O

Split Collective I/O pp. 681-689

HDF5

MPIO Read <https://github.com/HDFGroup/hdf5/blob/868f23f94effea398e145fb92c647193eb65dbd7/src/H5FDmpio.c#L1217>

MPIO Write <https://github.com/HDFGroup/hdf5/blob/868f23f94effea398e145fb92c647193eb65dbd7/src/H5FDmpio.c#L1467>

Suggested Change

Deprecate setting a file view on a file

Possibly with a new File Open flag? “MPI_MODE_NO_VIEW”?

Why? Requiring a file view imposes a collective requirement on non-contiguous independent I/O

Introduce new “viewless” independent & collective I/O operations

MPI_FILE_NEW_READ_AT(fh, offset, file_datatype, file_count, mem_datatype, mem_count, buf, status)

MPI_FILE_NEW_READ_AT_ALL(fh, offset, file_datatype, file_count, mem_datatype, mem_count, buf, status)

MPI_FILE_NEW_WRITE_AT(fh, offset, file_datatype, file_count, mem_datatype, mem_count, buf, status)

MPI_FILE_NEW_WRITE_AT_ALL(fh, offset, file_datatype, file_count, mem_datatype, mem_count, buf, status)

Possibly add offset in buffer, to make symmetric?

MPI_FILE_NEW_READ_AT(fh, offset, file_datatype, file_count, mem_datatype, mem_count, buf, buf_off, status)

Better aligns w/object access?

And corresponding nonblocking routines

Support split collective?

Alternatives:

Local file views Avoids collective requirement

Possibly solved w/local file views